

**SYSTEM AND METHOD FOR POWER MANAGEMENT
IN A SERVER SYSTEM**

The Field of the Invention

5 The present invention relates to server systems. More particularly, the invention relates to a system and method for power management in a cPCI server system.

Background of the Invention

10 Management of power in server systems is important to avoid configurations that could cause cards or the system to fail. A series of power hungry cards can cause the power supply to exceed its rated capacity, forcing it to shutdown and bring down the entire server chassis. Existing products typically impose a maximum anticipated power limit that applies to all card slots, even if there is a range of useful applications that do not approve of this maximum. Extra power capacity is forced into the supply and hence added cost or undue current limitations on configurations typically exists.

15 It would be desirable for a server system to use a calculated power management method on a module-by-module basis, allowing the server system to be optimized for modules of differing power requirements.

Summary of the Invention

20 One form of the present invention provides a server system including a plurality of subsystems, each including an associated memory for storing power usage information. A power supply unit coupled to the plurality of subsystems provides power to the plurality of subsystems. A server management card is coupled to the plurality of subsystems. The server management card is configured to retrieve the power usage information from the memory of each subsystem. The server management card is configured to calculate the total power usage of the plurality of subsystems based on the retrieved power usage information.

Brief Description of the Drawings

Figure 1 is a front perspective view illustrating a server system according to one embodiment of the present invention.

5 Figure 2 is a rear perspective view illustrating the server system shown in Figure 1.

Figure 3 is a block diagram illustrating major components of a server system according to one embodiment of the present invention.

10 Figure 4 is a front view of one of LCD panels used by a server system according to one embodiment of the present invention.

Figure 5 is an electrical block diagram illustrating major components of a server management card (SMC) according to one embodiment of the present invention.

Description of the Preferred Embodiments

15 In the following detailed description of the preferred embodiments, reference is made to the accompanying drawings that form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. It is to be understood that other embodiments may be utilized and structural or logical changes may be made without departing from the scope of the present invention. The following detailed description, therefore, is not to be taken in a limiting sense, and the scope of the present invention is defined by the appended claims.

I. **SERVER SYSTEM**

25 Figure 1 is a front perspective view illustrating a server system 100 according to one embodiment of the present invention. Figure 2 is a rear perspective view illustrating server system 100. Server system 100 includes panels 102, liquid crystal display (LCD) panels 104A and 104B (collectively referred to as LCD panels 104), backplane 106, chassis 108, and dual redundant power supply units 114A and 114B (collectively referred to as power supply units 114). Panels 102 are attached to chassis 108, and provide protection for the internal components of server system 100. Backplane 106 is positioned

near the center of server system 100. Backplane 106 is also referred to as midplane 106. LCD panels 104A and 104B are substantially identical, except for their placement on server system 100. LCD panel 104A is positioned on a front side of server system 100, and LCD panel 104B is positioned on a back side of server system 100.

5 Power supply units 114 are positioned at the bottom of server system 100 and extend from a back side of server system 100 to a front side of server system 100. Power supply units 114 each include an associated cooling fan 304 (shown in block form in Figure 3). In one form of the invention, additional cooling fans 304 are positioned behind LCD panel 104B. In one embodiment, 4 chassis cooling fans 304 are used in server 10 system 100. In an alternative embodiment, 6 chassis cooling fans 304 are used. Other numbers and placement of cooling fans 304 may be used. In one form of the invention, cooling fans 304 form an $N+1$ redundant cooling system, where “ N ” represents the total number of necessary fans 304, and “1” represents the number of redundant fans 304.

15 In one embodiment, server system 100 supports the Compact Peripheral Component Interconnect (cPCI) form factor of printed circuit assemblies (PCAs). Server system 100 includes a plurality of cPCI slots 110 for receiving cards/modules 300 (shown in block form in Figure 3). In one embodiment, system 100 includes ten slots 110 on each side of backplane 106 (referred to as the 10 slot configuration). In an alternative embodiment, system 100 includes nineteen slots 110 on each side of backplane 106 (referred to as the 19 slot configuration). Additional alternative embodiments use other 20 slot configurations.

Figure 3 is a block diagram illustrating major components of server system 100. Server system 100 includes backplane 106, a plurality of cards/modules 300A-300G (collectively referred to as cards 300), fans 304, electrically erasable programmable read 25 only memory (EEPROM) 314, LEDs 322, LCD panels 104, power supply units (PSUs) 114, and temperature sensor 324. Cards 300 are inserted in slots 110 (shown in Figures 1 and 2) in system 100. In one form of the invention, cards 300 may occupy more than one slot 110. In one embodiment, cards 300 include host processor cards 300A, hard disk cards 300B, managed Ethernet switch cards 300C and 300D, a server management card (SMC) 300E, and 2 redundant SMC local area network (LAN) rear transition modules 30

(RTMs) 300F and 300G. In one embodiment, there is one managed Ethernet switch card 300C fitted in the 10 slot chassis embodiment, and up to two managed Ethernet switch cards 300C and 300D fitted in the 19 slot chassis embodiment. In one form of the invention, managed Ethernet switch cards 300C and 300D are “Procurve” managed 5 Ethernet switch cards.

In one embodiment, two types of host processor cards 300A may be used in server system 100 – PA-RISC host processor cards and IA32 host processor cards. Multiple host processor cards 300A and hard disk cards 300B are used in embodiments of server system 100, but are each represented by a single card in Figure 3 to simply the figure. In 10 one form of the invention, up to 8 host processor cards 300A are used in the 10 slot configuration, and up to 16 host processor cards 300A are used in the 19 slot configuration. In one embodiment, each of cards 300 can be hot swapped.

In one embodiment, cards 300 each include a pair of EEPROMs 302A and 302B, which are discussed below. Power supply units 114 each include an EEPROM 323 for 15 storing power supply identification and status information. Fans 304 include associated sensors 306 for monitoring the speed of the fans 304. In one embodiment, LEDs 322 include eight status LEDs, six LAN LEDs to indicate the speed and link status of LAN links 318, a blue hot swap status LED to indicate the ability to hot swap SMC 300E, a power-on indicator LED, and three fan control indicator LEDs.

The operational health of cards 300 and system 100 are monitored by SMC 300E 20 to ensure the reliable operation of the system 100. SMC 300E includes serial ports 310 (discussed below), and an extraction lever 308 with an associated switch. In one embodiment, all cards 300 include an extraction lever 308 with an associated switch.

In one embodiment, SMC 300E is the size of a typical compact PCI (cPCI) card, 25 and supports PA-RISC and the IA32 host processor cards 300A. SMC 300E electrically connects to other components in system 100, including cards 300, temperature sensor 324, power supply units 114, fans 304, EEPROM 314, LCD panels 104, LEDs 322, and SMC rear transition modules 300F and 300G via backplane 106. In most cases, the connections are via I²C buses 554 (shown in Figure 5), as described in further detail 30 below. The I²C buses 554 allow bi-directional communication so that status information

can be sent to SMC 300E and configuration information sent from SMC 300E. In one embodiment, SMC 300E uses I²C buses 554 to obtain environmental information from power supply units 114, host processor cards 300A, and other cards 300 fitted into system 100.

5 SMC 300E also includes a LAN switch 532 (shown in Figure 5) to connect console management LAN signals from the host processor cards 300A to an external management network (also referred to as management LAN) 320 via one of the two SMC rear transition modules 300F and 300G. In one embodiment, the two SMC rear transition modules 300F and 300G each provide external 10/100Base-T LAN links 318 for
10 connectivity to management LAN 320. In one embodiment, SMC rear transition modules 300F and 300G are fibre-channel, port-bypass cards

Managed Ethernet switch cards 300C and 300D are connected to host processor cards 300A through backplane 106, and include external 10/100/1000Base-T LAN links 301 for connecting host processor cards to external customer or payload LANs 303.

15 15 Managed Ethernet switch cards 300C and 300D are fully managed LAN switches.

II. LCD PANEL

Figure 4 is a front view of one of LCD panels 104. In one form of the invention, each LCD panel 104 includes a 2 x 20 LCD display 400, 10 alphanumeric keys 402, 5 menu navigation/activation keys 404A-404E (collectively referred to as navigation keys 404), and a lockout key 406 with associated LED (not shown) that lights lockout key 406. If a user presses a key 402, 404, or 406, an alert signal is generated and SMC 300E polls the LCD panels 104A and 104B to determine which LCD panel was used, and the key that was pressed.

25 Alphanumeric keys 402 allow a user to enter alphanumeric strings that are sent to SMC 300E. Navigation keys 404 allow a user to navigate through menus displayed on LCD display 400, and select desired menu items. Navigation keys 404A and 404B are used to move left and right, respectively, within the alphanumeric strings. Navigation key 404C is an “OK/Enter” key. Navigation key 404D is used to move down. Navigation key 404E is a “Cancel” key.

LCD panels 104 provide access to a test shell (discussed below) that provides system information and allows configuration of system 100. As discussed below, other methods of access to the test shell are also provided by system 100. To avoid contention problems between the two LCD panels 104, and the other methods of access to the test 5 shell, a lockout key 406 is provided on LCD panels 104. A user can press lockout key 406 to gain or release control of the test shell. In one embodiment, lockout key 406 includes an associated LED to light lockout key 406 and indicate a current lockout status.

In one embodiment, LCD panels 104 also provide additional information to that displayed by LEDs 322 during start-up. If errors are encountered during the start-up 10 sequence, LCD panels 104 provide more information about the error without the operator having to attach a terminal to one of the SMC serial ports 310.

III. SERVER MANAGEMENT CARD (SMC)

A. SMC Overview

Figure 5 is an electrical block diagram illustrating major components of server management card (SMC) 300E. SMC 300E includes flash memory 500, processor 502, dynamic random access memory (DRAM) 504, PCI bridge 506, field programmable gate array (FPGA) 508, output registers 510A and 510B, input registers 512A and 512B, fan 20 controllers 526A-526C (collectively referred to as fan controllers 526), network controller 530, LAN switch 532, universal asynchronous receiver transmitter (UART) with modem 534, dual UART 536, UART with modem 538, clock generator/watchdog 540, battery 542, real time clock (RTC) 544, non-volatile random access memory (NVRAM) 546, I²C controllers 548A-548H (collectively referred to as I²C controllers 548), EEPROM 550, 25 and temperature sensor 324. In one embodiment, components of SMC 300E are connected together via PCI buses 507. In one form of the invention, PCI buses 507 are not routed between slots 110. Switched LAN signals through LAN switch 532 are routed between slots 110.

Functions of SMC 300E include supervising the operation of other components 30 within system 100 (e.g. fan speed, temperature, card present) and reporting their health to

a central location (e.g., external management network 320), reporting any failures to a central location (e.g., external management network 320), providing a LAN switch 532 to connect console management LAN signals from the SMC 300E and host processor cards 300A to an external management network 320, and providing an initial boot
5 configuration for the system 100.

B. SMC Processor And Memory

SMC 300E includes chassis management processor 502. In one embodiment, chassis management processor 502, also referred to as SMC processor 502, is a
10 StrongARM SA-110 processor with supporting buffer. In one embodiment, SMC 300E uses a Linux operating system. SMC 300E also runs server management application (SMA) software/firmware. In one embodiment, the operating system and SMA are stored in flash memory 500. In one form of the invention, all information needed to power-up SMC 300E, and for SMC 300E to become operational, are stored in flash memory 500.
15 In one embodiment, flash memory 500 includes 4 to 16 Mbytes of storage space to allow SMC 300E to boot-up as a stand-alone card (i.e., no network connection needed).

SMC 300E also includes DRAM 504. In one embodiment, DRAM 504 includes 32, 64 or 128 Mbytes of storage space. In one form of the invention, a hardware fitted table is stored in DRAM 504. The hardware fitted table includes information
20 representing the physical configuration of system 100. The hardware fitted table changes if there is a physical change to system 100, such as by a hardware device being added to or removed from system 100. The hardware fitted table includes hardware type information (e.g., whether a device is an IA32 / PA-RISC / Disk Carrier / RTM (i.e., rear transition module) / PSU / LCD panel / Modem / Unknown device, etc.), hardware revision and serial number, status information, configuration information, and hot-swap
25 status information.

Processor 502 is coupled to FPGA 508. FPGA 508 includes 6 sets of input/output lines 522A-522F. Lines 522A are connected to jumpers for configuring SMC 300E. Lines 522B are hot swap lines for monitoring the hot swap status of cards 300. In one embodiment, hot swap lines 522B include 18 hot swap status input lines, which allow
30

SMC 300E to determine the hot swap status of the host processor cards 300A, hard disk cards 300B, managed Ethernet switch cards 300C and 300D, SMC rear transition modules 300F and 300G, and power supply units 114. Lines 522C are LED lines that are coupled to LEDs 322. Lines 522D are fan input lines that are coupled to fan sensors 306 5 for monitoring the speed of fans 304. Lines 522E are power supply status lines that are coupled to power supply units 114 for determining whether both, or only one power supply unit 114 is present. Lines 522F are SMB alert lines for communicating alert signals related to SMB I²C buses 554B, 554D, and 554F.

10 C. Clock, Battery & NVRAM

SMC 300E includes a real time clock (RTC) 544 and an associated battery 542 to preserve the clock. Real time clock 544 provides the correct time of day. SMC 300E also includes NVRAM 546 for storing clock information. In one embodiment, NVRAM 546 uses the same battery as real time clock 544.

15 D. LAN switch

SMC 300E sends and receives management LAN communications through PCI bridge 506 and controller 530 to LAN switch 532. In one embodiment, LAN switch 532 is an unmanaged LAN switch including 19 ports, with two ports connected to SMC rear transition modules 300F and 300G (shown in Figure 3) via links 531A for communications with external management network 320 (shown in Figure 3), 16 ports for connecting to the management LAN connections of up to 16 host processor cards 300A via links 531B through backplane 106, and one port for connecting to the SMC's LAN port (i.e., output of controller 530) via links 531C. SMC 300E provides 20 management support for console LAN management signals sent and received through LAN switch 532. SMC 300E provides control of management LAN signals of host processor cards 300A, managed Ethernet switches 300C and 300D, SMC processor 502, and SMC rear transition modules 300F and 300G. SMC 300E monitors the status of the management LAN connections of up to 16 host processor cards 300A to LAN switch 532, 25 and reports an alarm event if any of the connections are lost. FPGA 508 and LAN switch 30

532 are coupled together via an RS-232 link 533 for the exchange of control and status information.

E. I²C Buses

5 Server system 100 includes eight I²C buses 554A-554H (collectively referred to as I²C buses 554) to allow communication with components within system 100. I²C buses 554 are coupled to FPGA 508 via I²C controllers 548. In one embodiment, the I²C buses 554 include 3 intelligent platform management bus (IPMB) buses 554A, 554C, and 554E, 10 3 system management bus (SMB) buses 554B, 554D, and 554F, a backplane ID bus (BP) 554G, and an I²C bus 554H for accessing SMC EEPROM 550 and chassis temperature sensor 324. A different number and configuration of I²C buses 554 may be used depending upon the desired implementation. SMC 300E maintains a system event log (SEL) within non-volatile flash memory 500 for storing information gathered over I²C buses 554.

15 The IPMB I²C buses 554A, 554C, and 554E implement the intelligent platform management interface (IPMI) specification. The IPMI specification is a standard defining an abstracted interface to platform management hardware. IPMI is layered over the standard I²C protocol. SMC 300E uses one or more of the IPMB I²C buses 554A, 554C, and 554E to retrieve static data from each of the host processor cards 300A and hard disk 20 cards 300B. The static data includes identification information for identifying each of the cards 300A and 300B. Each slot 110 in system 100 can be individually addressed to retrieve the static configuration data for the card 300 in that slot 110. In one embodiment, the host processor cards 300A and hard disk cards 300B each include an EEPROM 302A (shown in Figure 3) that stores the static identification information retrieved over IPMB 25 I²C buses 554A, 554C, and 554E. In one embodiment, each EEPROM 302A contains the type of card, the name of the card, the hardware revision of the card, the card's serial number and card manufacturing information.

SMC 300E also uses one or more of the IPMB I²C buses 554A, 554C, and 554E, to retrieve dynamic environmental information from each of the host processor cards 300A and hard disk cards 300B. In one embodiment, this dynamic information is held in

a second EEPROM 302B (shown in Figure 3) on each of the cards 300A and 300B. In one form of the invention, the dynamic board data includes card temperature and voltage measurements. In one embodiment, SMC 300E can write information to the EEPROMs 302A and 302B on cards 300.

5 The three SMB I²C buses 554B, 554D, and 554F also implement the IPMI specification. The three SMB I²C buses 554B, 554D, and 554F, are coupled to LEDs 322, the two LCD panels 104, the dual redundant power supply units 114, and some of the host processor cards 300A. SMC 300E uses one or more of the SMB I²C buses 554B, 554D, and 554F, to provide console communications via the LCD panels 104. In order
10 for the keypad key-presses on the LCD panels 104 to be communicated back to SMC 300E, an alert signal is provided when keys are pressed that causes SMC 300E to query LCD panels 104 for the keys that were pressed.

15 SMC 300E communicates with power supply units 114 via one or more of the SMB I²C buses 554B, 554D, and 554F to obtain configuration and status information including the operational state of the power supply units 114. In one embodiment, the dual redundant power supply units 114 provide voltage rail measurements to SMC 300E. A minimum and maximum voltage value is stored by the power supply units 114 for each measured rail. The voltage values are polled by SMC 300E at a time interval defined by the current configuration information for SMC 300E. If a voltage measurement goes out
20 of specification, defined by maximum and minimum voltage configuration parameters, SMC 300E generates an alarm event. In one embodiment, power supply units 114 store configuration and status information in their associated EEPROMs 323 (shown in Figure 3).

25 Backplane ID Bus (BP) 554G is coupled to backplane EEPROM 314 (shown in Figure 3) on backplane 106. SMC 300E communicates with the backplane EEPROM 314 over the BP bus 554G to obtain backplane manufacturing data, including hardware identification and revision number. On start-up, SMC 300E communicates with EEPROM 314 to obtain the manufacturing data, which is then added to the hardware fitted table. The manufacturing data allows SMC 300E to determine if it is in the correct
30 chassis for the configuration it has on board, since it is possible that the SMC 300E has

been taken from a different chassis and either hot-swapped into a new chassis, or added to a new chassis and the chassis is then powered up. If there is no valid configuration on board, or SMC 300E cannot determine which chassis it is in, then SMC 300E waits for a pushed configuration from external management network 320, or for a manual user configuration via one of the connection methods discussed below.

5 In one embodiment, there is a single temperature sensor 324 within system 100. SMC 300E receives temperature information from temperature sensor 324 over I²C bus 554H. SMC 300E monitors and records this temperature and adjusts the speed of the cooling fans 304 accordingly, as described below. SMC also uses I²C bus 554H to access 10 EEPROM 550, which stores board revision and manufacture data for SMC 300E.

F. Serial Ports

SMC 300E includes 4 RS-232 interfaces 310A-310D (collectively referred to as 15 serial ports 310). RS-232 serial interface 310A is via a 9-pin Male D-type connector on the front panel of SMC 300E. The other three serial ports 310B-310D are routed through backplane 106. The front panel RS-232 serial interface 310A is connected via a UART with a full modem 534 to FPGA 508, to allow monitor and debug information to be made 20 available via the front panel of SMC 300E. Backplane serial port 310D is also connected via a UART with a full modem 538 to FPGA 508. In one embodiment, backplane serial port 310D is intended as a debug or console port. The other two backplane serial interfaces 310B and 310C are connected via a dual UART 536 to FPGA 508, and are 25 routed to managed Ethernet switches 300C and 300D through backplane 106. These two backplane serial interfaces 310B and 310C are used to connect to and configure the managed Ethernet switch cards 300C and 300D, and to obtain status information from the managed Ethernet switch cards 300C and 300D.

G. Fans And Temperature Control

In one embodiment, server system 100 includes six chassis fans 304. Server 30 system 100 includes temperature sensor 324 to monitor the chassis temperature, and fan sensors 306 to monitor the six fans 304. In one embodiment, fan sensors 306 indicate

whether a fan 304 is rotating and the fan's speed setting. In one form of the invention, FPGA 508 includes 6 fan input lines 522D (i.e., one fan input line 522D from each fan sensor 306) to monitor the rotation of the six fans 304, and a single fan output line 524 coupled to fan controllers 526A-526C. Fan controllers 526A-526C control the speed of fans 304 by a PWM (pulse width modulation) signal via output lines 528A-528F. If a fan 304 stalls, the monitor line 522D of that fan 304 indicates this condition to FPGA 508, and an alarm event is generated. The speed of fans 304 is varied to maintain an optimum operating temperature versus fan noise within system 100. If the chassis temperature sensed by temperature sensor 324 reaches or exceeds a temperature alarm threshold, an alarm event is generated. When the temperature reduces below the alarm threshold, the alarm event is cleared. If the temperature reaches or exceeds a temperature critical threshold, the physical integrity of the components within system 100 are considered to be at risk, and SMC 300E performs a system shut-down, and all cards 300 are powered down except SMC 300E. When the chassis temperature falls below the critical threshold and has reached the alarm threshold, SMC 300E restores the power to all of the cards 300 that were powered down when the critical threshold was reached.

In one embodiment, SMC 300E controls the power state of cards 300 using power reset (PRST) lines 514 and power off (PWR_OFF) lines 516. FPGA 508 is coupled to power reset lines 514 and power off lines 516 via output registers 510A and 510B, respectively. In one embodiment, power reset lines 514 and power off lines 516 each include 19 output lines that are coupled to cards 300. SMC 300E uses power off lines 516 to turn off the power to selected cards 300, and uses power reset lines 514 to reset selected cards 300. In one embodiment, a lesser number of power reset and power off lines are used for the 10 slot chassis configuration.

25

H. Clock Generator / Watchdog

SMC 300E is protected by both software and hardware watchdog timers. The watchdog timers are part of clock generator/watchdog block 540, which also provides a clock signal for SMC 300E. The hardware watchdog timer is started before software loading commences to protect against failure. In one embodiment, the time interval is set

long enough to allow a worst-case load to complete. If the hardware watchdog timer expires, SMC processor 502 is reset.

5 I. Modes Of Operation

In one embodiment, SMC 300E has three phases or modes of operation -- Start-up, normal operation, and hot swap. The start-up mode is entered on power-up or reset, and controls the sequence needed to make SMC 300E operational. SMC 300E also provides minimal configuration information to allow chassis components to communicate on the management LAN. The progress of the start-up procedure can be followed on

10 LEDs 322, which also indicate any errors during start-up.

The normal operation mode is entered after the start-up mode has completed. In the normal operation mode, SMC 300E monitors the health of system 100 and its components, and reports alarm events. SMC 300E monitors the chassis environment, including temperature, fans, input signals, and the operational state of the host processor cards 300A.

15 SMC 300E reports alarm events to a central point, namely an alarm event manager, via the management LAN (i.e., through LAN switch 532 and one of the two SMC rear transition modules 300F or 300G to external management network 320). The alarm event manager is an external module that is part of external management network 320, and that handles the alarm events generated by server system 100. The alarm event manager decides what to do with received alarms and events, and initiates any recovery or reconfiguration that may be needed. In addition to sending the alarm events across the management network, a system event log (SEL) is maintained in SMC 300E to keep a record of the alarms and events. The SEL is held in non-volatile flash memory 500 in

20 SMC 300E and is maintained over power cycles, and resets of SMC 300E.

25 In the normal operation mode, SMC 300E may receive and initiate configuration commands and take action on received commands. The configuration commands allow the firmware of SMC processor 502 and the hardware controlled by processor 502 to be configured. This allows the operation of SMC 300E to be customized to the current environment. Configuration commands may originate from the management network

320, one of the local serial ports 310 via a test shell (discussed below), or one of the LCD panels 104.

The hot swap mode is entered when there is an attempt to remove a card 300 from system 100. In one embodiment, all of the chassis cards 300 can be hot swapped, including SMC 300E, and the two power supply units 114. An application shutdown sequence is initiated if a card 300 is to be removed. The shutdown sequence performs all of the steps needed to ready the card 300 for removal.

In one embodiment, FPGA 508 includes 18 hot swap status inputs 522B. These inputs 522B allow SMC 300E to determine the hot swap status of host processor cards 300A, hard disk cards 300B, managed Ethernet switch cards 300C and 300D, SMC rear transition module cards 300F and 300G, and power supply units 114. The hot-swap status of the SMC card 300E itself is also determined through this interface 522B.

An interrupt is generated and passed to SMC processor 502 if any of the cards 300 in system 100 are being removed or installed. SMC 300E monitors board select (BD_SEL) lines 518 and board healthy (HEALTHY) lines 520 of cards 300 in system 100. In one embodiment, board select lines 518 and healthy lines 520 each include 19 input lines, which are connected to FPGA 508 via input registers 512A and 512B, respectively. SMC 300E monitors the board select lines 518 to sense when a card 300 is installed. SMC 300E monitors the healthy lines 520 to determine whether cards 300 are healthy and capable of being brought out of a reset state.

When SMC 300E detects that a card has been inserted or removed, an alarm event is generated. When a new card 300 is inserted in system 100, SMC 300E determines the type of card 300 that was inserted by polling the identification EEPROM 302A of the card 300. Information is retrieved from the EEPROM 302A and added to the hardware fitted table. SMC 300E also configures the new card 300 if it has not been configured, or if its configuration differs from the expected configuration. When a card 300, other than the SMC 300E, is hot-swapped out of system 100, SMC 300E updates the hardware fitted table accordingly.

In one embodiment, SMC 300E is extracted in three stages: (1) an interrupt is generated and passed to the SMC processor 502 when the extraction lever 308 on the

SMC front panel is set to the “extraction” position in accordance with the Compact PCI specification, indicating that SMC 300E is about to be removed; (2) SMC processor 502 warns the external management network 320 of the SMC 300E removal and makes the extraction safe; and (3) SMC processor 502 indicates that SMC may be removed via the 5 blue hot swap LED 322. SMC 300E ensures that any application download and flashing operations are complete before the hot swap LED 322 indicates that the card 300E may be removed.

10 **J. User Connectivity**

In one embodiment, there are two test shells implemented within SMC 300E.

There is an application level test shell that is a normal, run-time, test shell accessed and used by users and applications. There is also a stand-alone test shell that is a manufacturer test shell residing in flash memory 500 that provides manufacturing level diagnostics and functions. The stand-alone test shell is activated when SMC 300E boots 15 and an appropriate jumper is in place on SMC 300E. The stand-alone test shell allows access to commands that the user would not, or should not have access to.

The test shells provide an operator interface to SMC 300E. This allows an operator to query the status of system 100 and (with the required authority level) to change the configuration of system 100.

20 A user can interact with the test shells by a number of different methods, including locally via a terminal directly attached to one of the serial ports 310, locally via a terminal attached by a modem to one of the serial ports 310, locally via one of the two LCD panels 104, and remotely via a telnet session established through the management LAN 320. A user may connect to the test shells by connecting a terminal to either the front panel serial port 310A or rear panel serial ports 310B-310D of SMC 300E, 25 depending on the console/modem serial port configuration. The RS-232 and LAN connections provide a telnet console interface. LCD panels 104 provide the same command features as the telnet console interface. SMC 300E can function as either a dial-in facility, where a user may establish a link by calling to the modem, or as a dial-out facility, where SMC 300E can dial out to a configured number.

The test shells provide direct access to alarm and event status information. In addition, the test shells provides the user with access to other information, including temperature logs, voltage logs, chassis card fitted table, and the current setting of all the configuration parameters. The configuration of SMC 300E may be changed via the test shells. Any change in configuration is communicated to the relevant cards 300 in system 100. In one embodiment, configuration information downloaded via a test shell includes a list of the cards 300 expected to be present in system 100, and configuration data for these cards 300. The configuration information is stored in flash memory 500, and is used every time SMC 300E is powered up.

10

K. Power Management

In one embodiment, power usage values by watt are embedded in an identification (ID) EEPROM of each field replaceable unit (FRU), which includes cards 300 and fans 304. For example, cards 300 (shown in Figure 3) each include ID EEPROM 302A, and SMC 300E includes EEPROM 550, for storing power usage values of each of these cards. In one embodiment, fans 304 also include an ID EEPROM. In one form of the invention, the power rating of each FRU 300 and 304 is also visibly color-coded on the FRU's bulkhead or an appropriately placed label.

20 In one embodiment of the present invention, SMC 300E polls the ID EEPROMs 302A and 550 of the FRUs 300 and 304 via one of the I²C buses 554 to obtain the power usage of each FRU 300 and 304. SMC 300E also polls EEPROM 323 of power supply units 114, which stores the power capacity of the power supply units 114. SMC 300E compares the power usage values obtained from the FRU ID EEPROMs, with the overall power available in server system 100 obtained from the power supply unit's ID EEPROM 323, and determines if there is sufficient capacity to power up the FRUs 300 and 304. SMC 300E controls the power state of FRUs 300 and 304 based on the comparison of the power usage values with the overall power available. If there is not sufficient capacity to power up the FRUs 300 and 304, SMC 300E does not power up all FRUs 300 and 304, or 25 does not power up selected ones of the FRUs 300 and 304.

15

20

25

30

In one embodiment, there are a total of five voltage rails from power supply units 114, with each rail having a different capacity for power. In one form of the invention, a power supply unit 114 has maximum ratings of: $48V \times 2.5A = 120W$, $12V \times 24.0A = 288W$, $5V \times 120.0A = 600W$, $3.3V \times 150.0A = 495W$, and $-12V \times 1.5A = 18W$; with an 5 additional maximum total power constraint for the supply of 1200W. When a card 300 is inserted into a slot 110 of server system 100, SMC 300E compares the power usage values of the other FRUs 300 and 304 to the total power budget of the power supply 114, and determines if the maximum values will be exceeded. If the maximum values will be exceeded, SMC 300E does not power on the inserted card 300, and responds with an 10 error message that is displayed on LCD panel 104.

Since server system 100 can be configured in a semi-infinite number of ways through the loading of its several slots 110, making a configuration chart is difficult for known released modules, and impossible for unknown future power hungry modules. Thus, by having a weighted number system that is automatically calculated by SMC 15 300E, configurations that would compromise the power integrity of the system 100 can be automatically avoided. Also, since the power supply units 114 can output their abilities (stored in EEPROM 323), upgrading to a higher current supply can be integrated without changing the code or documentation of SMC 300E.

Although specific embodiments have been illustrated and described herein for 20 purposes of description of the preferred embodiment, it will be appreciated by those of ordinary skill in the art that a wide variety of alternate and/or equivalent implementations may be substituted for the specific embodiments shown and described without departing from the scope of the present invention. Those with skill in the chemical, mechanical, electro-mechanical, electrical, and computer arts will readily appreciate that the present 25 invention may be implemented in a very wide variety of embodiments. This application is intended to cover any adaptations or variations of the preferred embodiments discussed herein. Therefore, it is manifestly intended that this invention be limited only by the claims and the equivalents thereof.